arXiv:1606.04495v1 [cs.DS] 14 Jun 2016

Range Majorities and Minorities in Arrays

DJAMAL BELAZZOUGUI, CERIST, Algeria
TRAVIS GAGIE, University of Helsinki, Finland
J. IAN MUNRO, University of Waterloo, Canada
GONZALO NAVARRO, University of Chile
YAKOV NEKRICH, University of Waterloo, Canada

Karpinski and Nekrich (2008) introduced the problem of parameterized range majority, which asks us to preprocess a string of length n such that, given the endpoints of a range, one can quickly find all the distinct elements whose relative frequencies in that range are more than a threshold τ. Subsequent authors have reduced their time and space bounds such that, when τ is fixed at preprocessing time, we need either $\mathcal{O}(n \lg(1/\tau))$ space and optimal $\mathcal{O}(1/\tau)$ query time or linear space and $\mathcal{O}((1/\tau) \lg \lg \sigma)$ query time, where σ is the alphabet size. In this paper we give the first linear-space solution with optimal $\mathcal{O}(1/\tau)$ query time, even with variable τ (i.e., specified with the query). For the case when σ is polynomial on the computer word size, our space is optimally compressed according to the symbol frequencies in the string. Otherwise, either the compressed space is increased by an arbitrarily small constant factor or the time rises to any function in $(1/\tau) \cdot \omega(1)$. We obtain the same results on the complementary problem of parameterized range minority introduced by Chan et al. (2015), who had achieved linear space and $\mathcal{O}(1/\tau)$ query time with variable τ.

Categories and Subject Descriptors: E.1 [**Data structures**]; E.4 [**Coding and information theory**]: Data compaction and compression

General Terms: Arrays, Range Queries

Additional Key Words and Phrases: Compressed data structures, range majority and minority

1. INTRODUCTION

Finding frequent elements in a dataset is a fundamental operation in data mining. The most frequent elements can be difficult to spot when all the elements have nearly equal frequencies. In some cases, however, we are interested in the most frequent elements only if they really are frequent. For example, Misra & Gries [1982] showed how, given a string and a threshold $0 < \tau \leq 1$, by scanning the string twice and using $\mathcal{O}(1/\tau)$ space we can find all the distinct elements whose relative frequencies exceed τ. These elements are called the τ-majorities of the string. If the element universe is $[1..\sigma]$, their algorithm can run in linear time and $\mathcal{O}(\sigma)$ space. Demaine *et al.* [2002] rediscovered the algorithm and deamortized the cost per element; Karp *et al.* [2003] rediscovered it again, obtaining $\mathcal{O}(1/\tau)$ space and linear randomized time. As

Funded in part by Academy of Finland grant 268324 and Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003 (Chile). An early partial version of this article appeared in *Proc. WADS 2013*.

Authors' addresses: D. Belazzougui, Research Center on Technical and Scientific Information, Algeria; T. Gagie, Department of Computer Science, University of Helsinki; I. Munro and Y. Nekrich, David Cheriton School of Computer Science, University of Waterloo, Canada; G. Navarro, Department of Computer Science, University of Chile.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© YYYY ACM 1549-6325/YYYY/01-ARTA $15.00
DOI:http://dx.doi.org/10.1145/0000000.0000000

Cormode & Muthukrishnan [2003] put it, "papers on frequent items are a frequent item!".

Krizanc *et al.* [2005] introduced the problem of preprocessing the string such that later, given the endpoints of a range, we can quickly return the mode of that range (i.e., the most frequent element). They gave two solutions, one of which takes $\mathcal{O}\big(n^{2-2\epsilon}\big)$ space for any fixed positive $\epsilon \leq 1/2$, and answers queries in $\mathcal{O}(n^{\epsilon} \lg \lg n)$ time; the other takes $\mathcal{O}\big(n^2 \lg \lg n / \lg n\big)$ space and answers queries in $\mathcal{O}(1)$ time. Petersen [2008] reduced Krizanc *et al.*'s first time bound to $\mathcal{O}(n^{\epsilon})$ for any fixed non-negative $\epsilon < 1/2$, and Petersen & Grabowski [2009] reduced the second space bound to $\mathcal{O}\big(n^2 \lg \lg n / \lg^2 n\big)$. Chan *et al.* [2014] gave an $\mathcal{O}(n)$ space solution that answers queries in $\mathcal{O}\Big(\sqrt{n / \lg n}\Big)$ time. They also gave evidence suggesting we cannot easily achieve query time substantially smaller than $\sqrt{n}$ using linear space; however, the best known lower bound [Greve *et al.* 2010] says only that we cannot achieve query time $o(\lg(n)/\lg(sw/n))$ using s words of w bits each. Because of the difficulty of supporting range mode queries, Bose *et al.* [2005] and Greve *et al.* [2010] considered the problem of approximate range mode, for which we are asked to return an element whose frequency is at least a constant fraction of the mode's frequency.

Karpinski & Nekrich [2008] took a different direction, analogous to Misra & Gries' approach, when they introduced the problem of preprocessing the string such that later, given the endpoints of a range, we can quickly return the τ-majorities of that range. We refer to this problem as parameterized range majority. Assuming τ is fixed when we are preprocessing the string, they showed how we can store the string in $\mathcal{O}(n(1/\tau))$ space and answer queries in $\mathcal{O}\big((1/\tau)(\lg \lg n)^2\big)$ time. They also gave bounds for dynamic and higher-dimensional versions. Durocher *et al.* [2013] independently posed the same problem and showed how we can store the string in $\mathcal{O}(n \lg(1/\tau + 1))$ space and answer queries in $\mathcal{O}(1/\tau)$ time. Notice that, because there can be up to $1/\tau$ distinct elements to return, this time bound is worst-case optimal. Gagie *et al.* [2011] showed how to store the string in compressed space — i.e., $\mathcal{O}(n(H+1))$ bits, where H is the entropy of the distribution of elements in the string — such that we can answer queries in $\mathcal{O}((1/\tau) \lg \lg n)$ time. Note that $H \leq \lg \sigma$, thus $\mathcal{O}(n(H+1))$ bits is $\mathcal{O}(n)$ space. They also showed how to handle a variable τ and still achieve optimal query time, at the cost of increasing the space bound by a $(\lg n)$-factor. That is, they gave a data structure that stores the string in $\mathcal{O}(n(H+1))$ words such that later, given the endpoints of a range and τ, we can return the τ-majorities of that range in $\mathcal{O}(1/\tau)$ time. Chan *et al.* [2015] gave another solution for variable τ, which also has $\mathcal{O}(1/\tau)$ query time and uses $\mathcal{O}(n \lg n)$ space. As far as we know, these are all the relevant bounds for Karpinski & Nekrich's original problem, both for fixed and variable τ; they are summarized in Table I together with our new results. Related work includes dynamic structures [Elmasry *et al.* 2011], approximate solutions [Lai *et al.* 2008; Wei & Yi 2011], and encodings that do not access the string at query time [Navarro & Thankachan 2016].

In this paper we first consider the complementary problem of parameterized range minority, which was introduced by Chan *et al.* [2015] (and then generalized to trees by Durocher *et al.* [2016]). For this problem we are asked to preprocess the string such that later, given the endpoints of a range, we can return (if one exists) a distinct element that occurs in that range but is not one of its τ-majorities. Such an element is called a τ-minority for the range. At first, finding a τ-minority might seem harder than finding a τ-majority because, for example, we are less likely to find a τ-minority by sampling. Nevertheless, Chan *et al.* gave an $\mathcal{O}(n)$ space solution with $\mathcal{O}(1/\tau)$ query time even for variable τ. In Section 3 we give two warm-up results, also for variable τ:

Table I. Results for the problem of parameterized range majority on a string of length n over an alphabet $[1..\sigma]$ with $\sigma \leq n$ in which the distribution of the elements has entropy $H \leq \lg \sigma$. Note that all the spaces given in bits are in $\mathcal{O}(n)$ words.

source	space	time	τ is
[Karpinski & Nekrich 2008]	$\mathcal{O}(n(1/\tau))$ words	$\mathcal{O}\big((1/\tau)(\lg\lg n)^2\big)$	fixed
[Durocher *et al.* 2013]	$\mathcal{O}(n\lg(1/\tau))$ words	$\mathcal{O}(1/\tau)$	fixed
[Gagie *et al.* 2011]	$\mathcal{O}(n(H+1))$ bits	$\mathcal{O}((1/\tau)\lg\lg\sigma)$	fixed
[Gagie *et al.* 2011]	$\mathcal{O}(n(H+1))$ words	$\mathcal{O}(1/\tau)$	variable
[Chan *et al.* 2015]	$\mathcal{O}(n\lg n)$ words	$\mathcal{O}(1/\tau)$	variable
Theorem 4.5 ($\lg\sigma = \mathcal{O}(\lg w)$)	$nH + o(n)$ bits	$\mathcal{O}(1/\tau)$	variable
Theorem 5.2	$nH + o(n)(H+1)$ bits	Any $(1/\tau)\cdot\omega(1)$	variable
Theorem 5.3	$(1+\epsilon)nH + o(n)$ bits	$\mathcal{O}(1/\tau)$	variable

(a) $\mathcal{O}(1/\tau)$ query time using $(1+\epsilon)nH + \mathcal{O}(n)$ bits of space for any constant $\epsilon > 0$; and
(b) any query time of the form $(1/\tau)\cdot\omega(1)$ using $nH + o(n)(H+1)$ bits of space.

That is, our spaces are not only linear like Chan *et al.*'s, but also compressed: we use either nearly optimally compressed space with no slowdown, or optimally compressed space with nearly no slowdown. Within this space, we can access S in constant time. We reuse ideas from this section in our solutions for parameterized range majority.

In Section 4 we return to Karpinski & Nekrich's original problem of parameterized range majority, but with variable τ, and consider the case where the alphabet of the string is polynomial on the computer word size, $\lg\sigma = \mathcal{O}(\lg w)$. We start with a simple linear-space structure (i.e., $\mathcal{O}(n\lg\sigma)$ bits, which is stricter than other linear-space solutions using $\mathcal{O}(n\lg n)$ bits) with worst-case optimal $\mathcal{O}(1/\tau)$ query time, and then refine it to obtain optimally compressed space. In Section 5 we extend this solution to the more challenging case of larger alphabets, $\lg\sigma = \omega(\lg w)$. We also start with a linear-space structure and then work towards compressing it. While we obtain nearly optimally compressed space and worst-case optimal query time, reaching optimally compressed space imposes a slight slowdown in the query time. Summarizing, we obtain the following tradeoffs:

(1) $\mathcal{O}(1/\tau)$ query time using $nH + o(n)$ bits of space, where $\lg\sigma = \mathcal{O}(\lg w)$;
(2) $\mathcal{O}(1/\tau)$ query time using $(1+\epsilon)nH + o(n)$ bits of space for any constant $\epsilon > 0$; and
(3) any query time of the form $(1/\tau)\cdot\omega(1)$ using $nH + o(n)(H+1)$ bits of space.

Overall, we obtain for the first time $\mathcal{O}(n)$-word (i.e., linear space) and $\mathcal{O}(1/\tau)$ (i.e., worst-case optimal) query time, even for variable τ. In all cases, we preserve constant-time access to S within the compressed space we use, and our queries require $\mathcal{O}(1/\tau)$ extra working space. As a byproduct, we also show how to find the range mode in a time that depends on how frequent it actually is.

Finally, in Section 6 we return to τ-minorities. By exploiting the duality with τ-majorities, we reuse the results obtained for the latter, so that the tradeoffs (1)–(3) are also obtained for τ-minorities. Those results supersede the original tradeoff (a) obtained in Section 3. Actually, a single data structure with the spaces given in points (1)–(3) solves at the same time τ-minority and τ-majority queries.

2. PRELIMINARIES

We use the RAM model of computation with word size in bits $w = \Omega(\lg n)$, allowing multiplications. The input is an array $S[1..n]$ of symbols (or "elements") from $[1..\sigma]$, where for simplicity we assume $\sigma \leq n$ (otherwise we could remap the alphabet so that

every symbol actually appears in S, without changing the output of any τ-majority or τ-minority query).

2.1. Access, select, rank, and partial rank

Let $S[1..n]$ be a string over alphabet $[1..\sigma]$, for $\sigma \leq n$, and let $H \leq \lg \sigma$ be the entropy of the distribution of elements in S, that is, $H = \sum_{a \in [1..\sigma]} \frac{n_a}{n} \lg \frac{n}{n_a}$, where each element a appears n_a times in S. An access query on S takes a position k and returns $S[k]$; a rank query takes an alphabet element a and a position k and returns $\mathrm{rank}_a(S, k)$, the number of occurrences of a in $S[1..k]$; a select query takes an element a and a rank r and returns $\mathrm{select}_a(S, r)$, the position of the rth occurrence of a in S. A partial rank query, $\mathrm{rank}_{S[k]}(S, k)$, is a rank query with the restriction that the element a must occur in the position k. These are among the most well-studied operations on strings, so we state here only the results most relevant to this paper.

For $\sigma = 2$ and any constant c, Pătraşcu [2008] showed how we can store S in $nH + \mathcal{O}(n/\lg^c n)$ bits, supporting all the queries in time $\mathcal{O}(c)$. If S has m 1s, this space is $\mathcal{O}\big(m \lg \frac{n}{m} + n/\lg^c n\big)$ bits. For $\lg \sigma = \mathcal{O}(\lg \lg n)$, Ferragina *et al.* [2007] showed how we can store S in $nH + o(n)$ bits and support all the queries in $\mathcal{O}(1)$ time. This result was later extended to the case $\lg \sigma = \mathcal{O}(\lg w)$ [Belazzougui & Navarro 2015, Thm. 7]. Barbay *et al.* [2014] showed how, for any positive constant ϵ, we can store S in $(1 + \epsilon)nH + o(n)$ bits and support access and select in $\mathcal{O}(1)$ time and rank in $\mathcal{O}(\lg \lg \sigma)$ time. Alternatively, they can store S in $nH + o(n)(H + 1)$ bits and support either access or select in time $\mathcal{O}(1)$, and the other operation, as well as rank, in time $\mathcal{O}(\lg \lg \sigma)$. Belazzougui & Navarro [2015, Thm. 8] improved the time of rank to $\mathcal{O}(\lg \lg_w \sigma)$, which they proved optimal, and the time of the non-constant operation to any desired function in $\omega(1)$. Belazzougui & Navarro [2014, Sec. 3] showed how to support $\mathcal{O}(1)$-time partial rank using $o(n)(H + 1)$ further bits. Throughout this article we also use some simpler variants of these results.

2.2. Colored range listing

Muthukrishnan [2002] showed how we can store $S[1..n]$ such that, given the endpoints of a range, we can quickly list the distinct elements in that range and the positions of their leftmost occurrences therein. Let $C[1..n]$ be the array in which $C[k]$ is the position of the rightmost occurrence of the element $S[k]$ in $S[1..k-1]$ — i.e., the last occurrence before $S[k]$ itself — or 0 if there is no such occurrence. Notice $S[k]$ is the first occurrence of that distinct element in a range $S[i..j]$ if and only if $i \leq k \leq j$ and $C[k] < i$. We store C, implicitly or explicitly, and a data structure supporting $\mathcal{O}(1)$-time range-minimum queries on C that returns the position of the leftmost occurrence of the minimum in the range.

To list the distinct elements in a range $S[i..j]$, we find the position m of the leftmost occurrence of the minimum in the range $C[i..j]$; check whether $C[m] < i$; and, if so, output $S[m]$ and m and recurse on $C[i..m-1]$ and $C[m+1..j]$. This procedure is online — i.e., we can stop it early if we want only a certain number of distinct elements — and the time it takes per distinct element is $\mathcal{O}(1)$ plus the time to access C.

Suppose we already have data structures supporting access, select and partial rank queries on S, all in $\mathcal{O}(t)$ time. Notice $C[k] = \mathrm{select}_{S[k]}(S, \mathrm{rank}_{S[k]}(S, k) - 1)$, so we can also support access to C in $\mathcal{O}(t)$ time. Sadakane [2007] and Fischer [2010] gave $\mathcal{O}(n)$-bit data structures supporting $\mathcal{O}(1)$-time range-minimum queries. Therefore, we can implement Muthukrishnan's solution using $\mathcal{O}(n)$ extra bits such that it takes $\mathcal{O}(t)$ time per distinct element listed.

3. PARAMETERIZED RANGE MINORITY

Chan *et al.* [2015] gave a linear-space solution with $\mathcal{O}(1/\tau)$ query time for parameterized range minority, even for the case of variable τ (i.e., chosen at query time). They first build a list of $\lceil 1/\tau \rceil$ distinct elements that occur in the given range (or as many as there are, if fewer) and then check those elements' frequencies to see which are τ-minorities. There cannot be as many as $\lceil 1/\tau \rceil$ τ-majorities so, if there exists a τ-minority for that range, then at least one must be in the list. In this section we use a simple approach to implement this idea using compressed space; then we obtain more refined results in Section 6.

3.1. Compressed space

To support parameterized range minority on $S[1..n]$ in $\mathcal{O}(1/\tau)$ time, we store a data structure supporting $\mathcal{O}(1)$-time access, select and partial rank queries on S [Barbay *et al.* 2014] and a data structure supporting $\mathcal{O}(1)$-time range-minimum queries on the C array corresponding to S [Muthukrishnan 2002]. As seen in Section 2, for any positive constant ϵ, we can store these data structures in a total of $(1+\epsilon)nH + \mathcal{O}(n)$ bits. Given τ and endpoints i and j, in $\mathcal{O}(1/\tau)$ time we use Muthukrishnan's algorithm to build a list of $\lceil 1/\tau \rceil$ distinct elements that occur in $S[i..j]$ (or as many as there are, if fewer) and the positions of their leftmost occurrences therein. We check whether these distinct elements are τ-minorities using the following lemma:

LEMMA 3.1. *Suppose we know the position of the leftmost occurrence of an element in a range. Then we can check whether that element is a τ-minority or a τ-majority using a partial rank query and a select query on S.*

PROOF. Let k be the position of the first occurrence of a in $S[i..j]$. If $S[k]$ is the rth occurrence of a in S, then a is a τ-minority for $S[i..j]$ if and only if the $(r+\lfloor \tau(j-i+1) \rfloor)$th occurrence of a in S is strictly after $S[j]$; otherwise a is a τ-majority. That is, we can check whether a is a τ-minority for $S[i..j]$ by checking whether

$$\mathrm{select}_a(S, \mathrm{rank}_a(S,k) + \lfloor \tau(j-i+1) \rfloor) > j\,;$$

since $S[k] = a$, computing $\mathrm{rank}_a(S,k)$ is only a partial rank query. □

This gives us the following theorem, which improves Chan *et al.*'s solution to use nearly optimally compressed space with no slowdown.

THEOREM 3.2. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H. For any constant $\epsilon > 0$, we can store S in $(1+\epsilon)nH + \mathcal{O}(n)$ bits such that later, given the endpoints of a range and τ, we can return a τ-minority for that range (if one exists) in time $\mathcal{O}(1/\tau)$.*

3.2. Optimally compressed space

By changing our string representation to that of Belazzougui & Navarro [2015, Thm. 8], we can store our data structures for access, select and partial rank on S and range-minimum queries on C in a total of $nH + o(nH) + \mathcal{O}(n)$ bits at the cost of the select queries taking $\mathcal{O}(g(n))$ time, for any desired $g(n) = \omega(1)$; see again Section 2. Therefore the range minority is found in time $\mathcal{O}((1/\tau)g(n))$.

To reduce the space bound to $nH + o(n)(H+1)$ bits, we must reduce the space of the range-minimum data structure to $o(n)$. Such a result was sketched by Hon *et al.* [2009], but it lacks sufficient detail to ensure correctness. We give these details next.

The technique is based on sparsification. We cut the sequence into blocks of length $g(n)$, choose the $n/g(n)$ minimum values of each block, and build the range-minimum

data structure on the new array $C'[1..n/g(n)]$ (i.e., $C'[i]$ stores the minimum of $C[(i-1)\cdot g(n)+1..i\cdot g(n)]$). This requires $\mathcal{O}(n/g(n)) = o(n)$ bits. Muthukrishnan's algorithm is then run over C' as follows. We find the minimum position in C', then recursively process its left interval, then process the minimum of C' by considering the $g(n)$ corresponding cells in C, and finally process the right part of the interval. The recursion stops when the interval becomes empty or when all the $g(n)$ elements in the block of C are already reported.

LEMMA 3.3. *The procedure described identifies the leftmost positions of all the distinct elements in a block-aligned interval $S[i..j]$, working over at most $g(n)$ cells per new element discovered.*

PROOF. We proceed by induction on the size of the current subinterval $[\ell..r]$, which is always block-aligned. Let k' be the position of the minimum in $C'[(\ell-1)g(n)+1..r/g(n)]$ and let k be the position of the minimum in $C[(k'-1)\cdot g(n)+1, k'\cdot g(n)]$. Then $C[k]$ is clearly the minimum in $C[\ell..r]$ and $S[k]$ is the leftmost occurrence in $S[\ell..r]$ of the element $a = S[k]$. If $C[k] \geq i$, then a already occurs in $S[i..\ell-1]$ and we have already reported it. Since the minimum of $C[\ell..r]$ is within the block k' of C, it is sufficient that $C[k] \geq i$ for all the positions k in that block to ensure that all the values in $S[\ell..r]$ have already been reported, in which case we can stop the procedure. The $g(n)$ scanned cells can be charged to the function that recursively invoked the interval $[\ell..r]$.

Otherwise, we recursively process the interval to the left of block k', which by inductive hypothesis reports the unique elements in that interval. Then we process the current block of size $g(n)$, finding at least the new occurrence of element $S[k]$ (which cannot appear to the left of k'). Finally, we process the interval to the right of k', where the inductive hypothesis again holds.

Note that the method is also correct if, instead of checking whether all the elements in the block of $C'[k']$ are $\geq i$, we somehow check that all of them have already been reported. We will use this variant later in the paper. □

For general ranges $S[i..j]$, we must include in the range of C' the two partially overlapped blocks on the extremes of the range. When it comes to process one of those blocks, we only consider the cells that are inside $[i..j]$; the condition to report an element is still that $C[k] < i$.

We use this procedure to obtain any $\lceil 1/\tau \rceil$ distinct elements. We perform $\mathcal{O}((1/\tau)\, g(n))$ accesses to C, each of which costs time $\mathcal{O}(g(n))$ because it involves a select query on S. Therefore the total time is $\mathcal{O}\big((1/\tau)\, g(n)^2\big)$. Testing each of the candidates with Lemma 3.1 takes time $\mathcal{O}((1/\tau)\, g(n))$, because we also use select queries on S. Therefore, for any desired time of the form $\mathcal{O}((1/\tau)\, f(n))$, we use $g(n) = \sqrt{f(n)}$.

THEOREM 3.4. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H. For any function $f(n) = \omega(1)$, we can store S in $nH + o(n)(H+1)$ bits such that later, given the endpoints of a range and τ, we can return a τ-minority for that range (if one exists) in time $\mathcal{O}((1/\tau)\, f(n))$.*

Note that this representation retains constant-time access to S.

4. PARAMETERIZED RANGE MAJORITY ON SMALL ALPHABETS

In this section we consider the case $\lg \sigma = \mathcal{O}(\lg w)$, where rank queries on S can be supported in constant time. Our strategy is to find a set of $\mathcal{O}(1/\tau)$ candidates that contain all the possible τ-majorities and then check them one by one, counting their occurrences in $S[i..j]$ via rank queries on S. The time will be worst-case optimal, $\mathcal{O}(1/\tau)$. We first obtain $\mathcal{O}(n \lg \sigma)$ bits of space and then work towards compressing it.

First, note that if $\tau < 1/\sigma$, we can simply assume that all the σ symbols are candidates for majority, and check them one by one; therefore we care only about how to find $\mathcal{O}(1/\tau)$ candidates in the case $\tau \geq 1/\sigma$.

4.1. Structure

We store an instance of the structure of Belazzougui & Navarro [2015, Thm. 5] supporting access, rank, and select on S in $\mathcal{O}(1)$ time, using $n \lg \sigma + o(n)$ bits. For every $0 \leq t \leq \lceil \lg \sigma \rceil$ and $t \leq b \leq \lfloor \lg n \rfloor$, we divide S into blocks of length 2^{b-1} and store a binary string $G_b^t[1..n]$ in which $G_b^t[k] = 1$ if, (1) the element $S[k]$ occurs at least 2^{b-t} times in $S[k - 2^{b+1}..k + 2^{b+1}]$, and (2) $S[k]$ is the leftmost or rightmost occurrence of that element in its block.

At query time we will use $t = \lceil \lg(1/\tau) \rceil$ and $b = \lfloor \lg(j - i + 1) \rfloor$. The following lemma shows that it is sufficient to consider the candidates $S[k]$ for $i \leq k \leq j$ where $G_b^t[k] = 1$.

LEMMA 4.1. *For every τ-majority a of $S[i..j]$ there exists some $k \in [i..j]$ such that $S[k] = a$ and $G_b^t[k] = 1$.*

PROOF. Since $S[i..j]$ cannot be completely contained in a block of length 2^{b-1}, if $S[i..j]$ overlaps a block then it includes one of that block's endpoints. Therefore, if $S[i..j]$ contains an occurrence of an element a, then it includes the leftmost or rightmost occurrence of a in some block. Suppose a is a τ-majority in $S[i..j]$, and $b \geq t$. For all $i \leq k \leq j$, a occurs at least $\tau 2^b \geq 2^{b-t}$ times in $S[k - 2^{b+1}..k + 2^{b+1}]$, so since some occurrence of a in $S[i..j]$ is the leftmost or rightmost in its block, it is flagged by a 1 in $G_b^t[i..j]$. □

The number of distinct elements that occur at least 2^{b-t} times in a range of size $\mathcal{O}(2^b)$ is $\mathcal{O}(2^t)$, so in each block there are $\mathcal{O}(2^t)$ positions flagged by 1s in G_b^t, for a total of $m = \mathcal{O}(n\, 2^{t-b})$ 1s. It follows that we can store an instance of the structure of Pătraşcu [2008] (recall Section 2.1) supporting $\mathcal{O}(1)$-time access, rank and select on G_b^t in $\mathcal{O}(n2^{t-b}(b-t) + n/\lg^3 n)$ bits in total. Summing over t from 0 to $\lceil \lg \sigma \rceil$ and over b from t to $\lfloor \lg n \rfloor$, calculation shows we use a total of $\mathcal{O}(n \lg \sigma)$ bits for the binary strings.

4.2. Queries

Given endpoints i and j and a threshold τ, if $\tau < 1/\sigma$, we simply report every element $a \in [1..\sigma]$ such that $\mathrm{rank}_a(S, j) - \mathrm{rank}_a(S, i-1) > \tau(j-i+1)$, in total time $\mathcal{O}(\sigma) = \mathcal{O}(1/\tau)$. Otherwise, we compute b and t as explained and, if $b < t$, we run a sequential algorithm on $S[i..j]$ in $\mathcal{O}(j - i) = \mathcal{O}(1/\tau)$ time [Misra & Gries 1982]. Otherwise, we use rank and select on G_b^t to find all the 1s in $G_b^t[i..j]$. Since $S[i..j]$ overlaps at most 5 blocks of length 2^{b-1}, it contains $\mathcal{O}(1/\tau)$ elements flagged by 1s in G_b^t; therefore, we have $\mathcal{O}(1/\tau)$ candidates to evaluate, and these include all the possible τ-majorities. Each candidate a is tested in constant time for the condition $\mathrm{rank}_a(S, j) - \mathrm{rank}_a(S, i-1) > \tau(j - i + 1)$.

We aim to run the sequential algorithm in $\mathcal{O}(j - i) = \mathcal{O}(1/\tau)$ worst-case time and space, which we achieve by taking advantage of the rank and select operations on S. We create a doubly-linked list with the positions i to j, plus an array $T[1..j-i+1]$ where $T[k]$ points to the list node representing $S[i+k-1]$. We take the element $S[i] = a$ at the head of the list and know that it is a τ-majority in $S[i..j]$ if $\mathrm{rank}_a(S, j) - \mathrm{rank}_a(S, i-1) > \tau(j-i+1)$. If it is, we immediately report it. In any case, we remove all the occurrences of a from the doubly-linked list, that is, the list nodes $T[\mathrm{select}_a(S, \mathrm{rank}_a(S, i) + r)]$, $r = 0, 1, 2, \ldots$. We proceed with the new header of the doubly-linked list, which points to a different element $S[i'] = a'$, and so on. It is clear that we perform $\mathcal{O}(j - i)$ constant-time rank and select operations on S, and that at the end we have found all the τ-majorities.

THEOREM 4.2. *Let $S[1..n]$ be a string over alphabet $[1..\sigma]$, with $\lg\sigma = \mathcal{O}(\lg w)$. We can store S in $\mathcal{O}(n\lg\sigma)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

4.3. Succinct space

To reduce the space we will open the structure we are using to represent S [Belazzougui & Navarro 2015, Thm. 5]. This is a multiary wavelet tree: it cuts the alphabet range $[1..\sigma]$ into w^β contiguous subranges of about the same size, for some conveniently small constant $0 < \beta \leq 1/4$. The root node v of the wavelet tree stores the sequence $S_v[1..n]$ indicating the range to which each symbol of S belongs, roughly $S_v[i] = \lceil S[i]/(\sigma/w^\beta)\rceil$. This node has w^β children, where the pth child stores the subsequence of the symbols $S[i]$ such that $S_v[i] = p$. The alphabet of each child has been reduced to a range of size roughly σ/w^β. This range is split again into w^β subranges, creating w^β children for each child, and so on. The process is repeated recursively until the alphabet range is of size less than w^β. The wavelet tree has height $\lg_{w^\beta}\sigma = \frac{\lg\sigma}{\beta\lg w} = \mathcal{O}(1)$, and at each level the strings S_v stored add up to $n\lg(w^\beta) = \beta n\lg w$ bits, for a total of $n\lg\sigma$ bits of space. The other $o(n)$ bits are needed to provide constant-time rank and select support on the strings S_v.

We use this hierarchical structure to find the τ-majorities as follows. Assume we have the structures to find τ-majorities in any of the strings S_v associated with wavelet tree nodes v, in time $\mathcal{O}(1/\tau)$. Then, if a is a τ-majority in $S[i..j]$, the symbol $p = \lceil a/(\sigma/w^\beta)\rceil$ is also a τ-majority in $S_v[i..j]$, where v is the wavelet tree root. Therefore, we find in time $\mathcal{O}(1/\tau)$ the τ-majorities p in $S_v[i..j]$. We verify each such τ-majority p recursively in the pth child of v. In this child u, the range $S_v[i..j]$ is projected to $S_u[i_u..j_u] = S_u[\mathrm{rank}_p(S_v, i-1)+1, \mathrm{rank}_p(S_v, j)]$, and the corresponding threshold is $\tau_u = \tau(j-i+1)/(j_u-i_u+1) < 1$. This process continues recursively until we find the majorities in the leaf nodes, which correspond to actual symbols that can be reported as τ-majorities in $S[i..j]$.

The time to find the τ_u-majorities in each child u of the root v is $\mathcal{O}(1/\tau_u) = \mathcal{O}((j_u-i_u+1)/((j-i+1)\tau))$. Added over all the children u, this gives $\sum_u \mathcal{O}(1/\tau_u) = \sum_u \mathcal{O}((j_u-i_u+1)/((j-i+1)\tau)) = \mathcal{O}(1/\tau)$. Adding this over all the levels, we obtain $\mathcal{O}((1/\beta)(1/\tau)) = \mathcal{O}(1/\tau)$.

Finding τ'-majorities on tiny alphabets. The remaining problem is how to find τ'-majorities on an alphabet of size $\sigma' = w^\beta$, on each of the strings S_v of length n_v. We do almost as we did for Theorem 4.2, except that the range for b is slightly narrower: $\lfloor\lg(2^t\cdot w^\beta/4)\rfloor \leq b \leq \lfloor\lg n_v\rfloor$. Then calculation shows that the total space for the bitvectors G_b^t is $\mathcal{O}\left(\frac{n_v\lg\sigma'\lg w}{w^\beta} + \frac{n_v}{\lg n_v}\right) = o(n_v)$, so added over the whole wavelet tree is $o(n)$.

The price of using this higher lower bound for b is that it requires us to sequentially find τ'-majorities in time $\mathcal{O}(1/\tau')$ on ranges of length $\mathcal{O}\big((1/\tau')w^\beta\big)$. However, we can take advantage of the small alphabet. First, if $1/\tau' \geq \sigma'$, we just perform σ' pairs of constant-time rank queries on S_v. For $1/\tau' < \sigma'$, we will compute an array of σ' counters with the frequency of the symbols in the range, and then report those exceeding the threshold. The maximum size of the range is $(4/\tau')w^\beta/4 \leq \sigma' w^\beta = w^{2\beta}$, and thus $2\beta\lg w$ bits suffice to represent each counter. The σ' counters then require $2\beta w^\beta\lg w$ bits and can be maintained in a computer word (although we will store them somewhat spaced for technical reasons). We can read the elements in S_v by chunks of w^β symbols, and compute in constant time the corresponding counters for those symbols. Then we sum the current counters and the counters for the chunk, all in constant time because they are fields in a single computer word. The range is then processed in time $\mathcal{O}(1/\tau')$.

To compute the counters corresponding to w^β symbols, we extend the popcounting algorithm of Belazzougui & Navarro [2015, Sec. 4.1]. Assume we extract them from S_v and have them packed in the lowest $k\ell$ bits of a computer word X, where $k = w^\beta$ is the number of symbols and $\ell = \lg \sigma'$ the number of bits used per symbol. We first create σ' copies of the sequence at distance $2k\ell$ of each other: $X \leftarrow X \cdot (0^{2k\ell-1}1)^{\sigma'}$. In each copy we will count the occurrences of a different symbol. To have the ith copy count the occurrences of symbol i, for $0 \le i < \sigma'$, we perform

$$X \leftarrow X \text{ XOR } 0^{k\ell}((\sigma'-1)_\ell)^k \ldots 0^{k\ell}(2_\ell)^k\, 0^{k\ell}(1_\ell)^k\, 0^{k\ell}(0_\ell)^k,$$

where i_ℓ is number i written in ℓ bits. Thus in the ith copy the symbols equal to i become zero and the others nonzero. To set a 1 at the highest bit of the symbols equal to i in the ith copy, we do

$$X \leftarrow (Y - (X \text{ AND NOT } Y)) \text{ AND } Y \text{ AND NOT } X,$$

where $Y = (0^{k\ell}(10^{\ell-1})^k)^{\sigma'}$.[1] Now we add all the 1s in each copy with $X \leftarrow X \cdot 0^{k\ell(2\sigma'-1)}(0^{\ell-1}1)^k$. This spreads several sums across the $2k\ell$ bits of each copy, and in particular the kth sum adds up all the 1s of the copy. Each sum requires $\lg k$ bits, which is precisely the ℓ bits we have allocated per field. Finally, we isolate the desired counters using $X \leftarrow X$ AND $(0^{k\ell}1^\ell 0^{(k-1)\ell})^{\sigma'}$. The σ' counters are not contiguous in the computer word, but we still can afford to store them spaced: we use $2k\ell\sigma' = 2\beta w^{2\beta} \lg w$ bits, which since $\beta \le 1/4$, is always less than w.

The cumulative counters, as said, need $\lg(\sigma' w^\beta) = 2\ell$ bits. We will store them in a computer word C separated by $2k\ell$ bits so that we can directly add the resulting word X after processing a chunk of w^β symbols of the range in S_v: $C \leftarrow C + X$. If the last chunk is of length $l < w^\beta$, we complete it with zeros and then subtract those spurious $w^\beta - l$ occurrences from the first counter, $C \leftarrow C - (w^\beta - l) \cdot 2^{(k-1)\ell}$.

The last challenge is to output the counters that are at least $y = \lfloor \tau'(j - i + 1) \rfloor + 1$ after processing the range. We use

$$C \leftarrow C + (2^{2\ell} - y) \cdot (0^{k\ell+\ell-1}10^{(k-1)\ell})^{\sigma'}$$

so that the counters reaching y will overflow to the next bit. We isolate those overflow bits with $C \leftarrow C$ AND $(0^{(k-1)\ell-1}10^{(k+1)\ell})^{\sigma'}$, so that we have to report the ith symbol if and only if C AND $0^{(k(2\sigma'-2i+1)-1)\ell-1}10^{(k(2i-1)+1)\ell} \neq 0$. We repeatedly isolate the lowest bit of C with

$$D \leftarrow (C \text{ XOR } (C-1)) \text{ AND } (0^{(k-1)\ell-1}10^{(k+1)\ell})^{\sigma'},$$

and then remove it with $C \leftarrow C$ AND $(C-1)$, until $C = 0$. Once we have a position isolated in D, we find the position in constant time by using a monotone minimum perfect hash function over the set $\{2^{(k(2i-1)+1)\ell},\ 1 \le i \le \sigma'\}$, which uses $\mathcal{O}(\sigma' \lg w) = o(w)$ bits [Belazzougui *et al.* 2009]. Only one such data structure is needed for all the sequences, and it takes less space than a single systemwide pointer.

THEOREM 4.3. *Let $S[1..n]$ be a string over alphabet $[1..\sigma]$, with $\lg \sigma = \mathcal{O}(\lg w)$. We can store S in $n \lg \sigma + o(n)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

[1] This could have been simply $X \leftarrow (Y - X)$ AND Y if there was an unused highest bit set to zero in the fields of X. Instead, we have to use this more complex formula that first zeroes the highest bit of the fields and later considers them separately.

4.4. Optimally compressed space

One choice to compress the space is to use a compressed representation of the strings S_v [Ferragina *et al.* 2007]. This takes chunks of $c = (\lg n)/2$ bits and assigns them a code formed by a header of $\lg c$ bits and a variable-length remainder of at most c bits. For decoding a chunk in constant time, they use a directory of $\mathcal{O}(n_v \lg c/c)$ bits, plus a constant table of size $2^c = \mathcal{O}(\sqrt{n})$ that receives any encoded string and returns the original chunk. The compressed size of any string S_v with entropy H_v then becomes $n_v H_v + \mathcal{O}(n_v \lg \sigma' \lg\lg n/\lg n)$ bits, which added over the whole wavelet tree becomes $nH + \mathcal{O}\left(\frac{n \lg \sigma \lg\lg n}{\lg n}\right)$ bits. This can be used in replacement of the direct representation of sequences S_v in Theorem 4.3, since we only change the way a chunk of $\Theta(\lg n)$ bits is read from any S_v. Note that we read chunks of w^β symbols from S_v, which could be $\omega(\lg n)$ if n is very small. To avoid this problem, we apply this method only when $\lg \sigma = \mathcal{O}(\lg\lg n)$, as in this case we can use computer words of $w = \lg n$ bits.

COROLLARY 4.4. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$, with $\lg \sigma = \mathcal{O}(\lg\lg n)$. We can store S in $nH + o(n)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

For the case where $\lg \sigma = \omega(\lg\lg n)$ but still $\lg \sigma = \mathcal{O}(\lg w)$, we use another technique. We represent S using the optimally compressed structure of Barbay *et al.* [2014]. This structure separates the alphabet symbols into $\lg^2 n$ classes according to their frequencies. A sequence $K[1..n]$, where $K[i]$ is the class to which $S[i]$ is assigned, is represented using the structure of Corollary 4.4, which supports constant-time access, rank, and select, since the alphabet of K is of polylogarithmic size, and also τ-majority queries in time $\mathcal{O}(1/\tau)$. For each class c, a sequence $S_c[1..n_c]$ contains the subsequence of S of the symbols $S[i]$ where $K[i] = c$. We will represent the subsequences S_c using Theorem 4.3. Then the structure for K takes $nH_K + o(n)$ bits, where H_K is the entropy of the distribution of the symbols in K, and the structures for the strings S_c take $n_c \lg \sigma_c + o(n_c)$ bits, where S_c ranges over alphabet $[1..\sigma_c]$. Barbay *et al.* show that these space bounds add up to $nH + o(n)$ bits and that one can support access, rank and select on S via access, rank and select on K and some S_c.

Our strategy to solve a τ-majority query on $S[i..j]$ resembles the one used to prove Theorem 4.3. We first run a τ-majority query on string K. This will yield the at most $1/\tau$ classes of symbols that, together, occur more than $\tau(j - i + 1)$ times in $S[i..j]$. The classes excluded from this result cannot contain symbols that are τ-majorities. Now, for each included class c, we map the interval $S[i..j]$ to $S_c[i_c..j_c]$ in the subsequence of its class, where $i_c = \mathrm{rank}_c(K, i-1) + 1$ and $j_c = \mathrm{rank}_c(K, j)$, and then run a τ_c-majority query on $S_c[i_c..j_c]$, for $\tau_c = \tau(j - i + 1)/(j_c - i_c + 1)$. The results obtained for each considered class c are reported as τ-majorities in $S[i..j]$. The query time, added over all the possible τ_c values, is $\sum_c \mathcal{O}(1/\tau_c) = \mathcal{O}(1/\tau)$ as before.

THEOREM 4.5. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$, with $\lg \sigma = \mathcal{O}(\lg w)$. We can store S in $nH + o(n)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

5. PARAMETERIZED RANGE MAJORITY ON LARGE ALPHABETS

Rank queries cannot be performed in constant time on large alphabets [Belazzougui & Navarro 2015]. To obtain optimal query time in this case, we resort to the use of Lemma 3.1 instead of performing rank queries on S. For this purpose, we must be able to find the leftmost occurrence of each τ-majority in a range. This is done

by adding further structures on top of the bitvectors G_b^t used in Theorem 4.2. Those bitvectors G_b^t alone require $\mathcal{O}(n \lg \sigma)$ bits of space, whereas our further structures add only $o(n)$ bits. Within these $\mathcal{O}(n \lg \sigma)$ bits, we can store a simple representation of S [Barbay *et al.* 2014], which supports both access and select queries in constant time. We also add the structures to support partial rank in constant time, within $o(n \lg \sigma)$ further bits. Therefore we can apply Lemma 3.1 in constant time and solve τ-majority queries in time $\mathcal{O}(1/\tau)$. We consider compression later.

5.1. Structure

First, to cover the case $\tau < 1/\sigma$, we build the structure of Muthukrishnan [2002] on S, using $\mathcal{O}(n)$ extra bits as shown in Section 2.2, so that we can find the $\mathcal{O}(\sigma) = \mathcal{O}(1/\tau)$ leftmost occurrences of each distinct element in $S[i..j]$. On each leftmost occurrence we can then apply Lemma 3.1 in constant time. Now we focus on the case $\tau \geq 1/\sigma$.

In addition to the bitvectors G_b^t of the previous section, we mark in a second bitvector J_b^t each $(\lg^4 n)$-th occurrence in S of the alphabet symbols in the area where they mark bits in G_b^t. More precisely, let $a = S[k]$ occur at least 2^{b-t} times in $S[k - 2^{b+1}..k + 2^{b+1}]$, and let $i_1, i_2, \ldots$ be the positions of a in S. Then we mark in J_b^t the positions $\{i_{q \lg^4 n}, k - 2^{b+1} \leq i_{q \lg^4 n} \leq k + 2^{b+1}\}$.

For the subsequence S_b^t of elements of S marked in J_b^t, we build an instance of Muthukrishnan's structure. That is, we build the structure on the array C_b^t corresponding to the string $S_b^t[k] = S[\mathrm{select}_1(J_b^t, k)]$. This string need not be stored explicitly, but instead we store C_b^t in explicit form.

Furthermore, if for any b and t it holds $J_b^t[i_{q \lg^4 n}] = 1$, being $S[i_{q \lg^4 n}] = a$, we create a succinct SB-tree [Grossi *et al.* 2009, Lem 3.3] successor structure[2] associated with the chunk of $\lg^4 n$ consecutive positions of a: $i_{1+(q-1) \lg^4 n}, \ldots, i_{q \lg^4 n}$. This structure is stored associated with the 1 at $J_b^t[i_{q \lg^4 n}]$ (all the 1s at the same position $i_{q \lg^4 n}$, for different b and t values, point to the same succinct SB-tree, as it does not depend on b or t). The SB-tree operates in time $\mathcal{O}\big(\lg(\lg^4 n)/\lg\lg n\big) = \mathcal{O}(1)$ and uses $\mathcal{O}\big(\lg^4 n \lg\lg n\big)$ bits. It needs constant-time access to the positions $i_{r+(q-1) \lg^4 n}$, as it does not store them. We provide those positions using $i_k = \mathrm{select}_a(S, k)$.

Added over all the symbols a, occurring n_a times in S, each bitvector J_b^t contains $\sum_a \lfloor n_a/\lg^4 n \rfloor = \mathcal{O}\big(n/\lg^4 n\big)$ 1s. Thus, added over every b and t, the bitvectors J_b^t, arrays C_b^t, and pointers to succinct SB-trees (using $\mathcal{O}(\lg n)$ bits per pointer), require $\mathcal{O}(n/\lg n) = o(n)$ bits. Each succinct SB-tree requires $\mathcal{O}\big(\lg^4 n \lg\lg n\big)$ bits, and they may be built for $\mathcal{O}\big(n/\lg^4 n\big)$ chunks, adding up to $\mathcal{O}(n \lg\lg n)$ bits. This is $\mathcal{O}(n \lg \sigma)$ if we assume $\lg \sigma = \Omega(\lg\lg n)$.

5.2. Queries

Given i and j, we compute $b = \lfloor \lg(j - i + 1) \rfloor$ and $t = \lceil \lg(1/\tau) \rceil$, and find the $\mathcal{O}(1)$ blocks of length 2^b overlapping $S[i..j]$. As in the previous section, every $G_b^t[k] = 1$ in $G_b^t[i..j]$ is a candidate to verify, but this time we need to find its leftmost occurrence in $S[i..j]$.

To find the leftmost position of $a = S[k]$, we see if the positions k and i are in the same chunk. That is, we compute the chunk index $q = \lceil \mathrm{rank}_a(S, k)/\lg^4 n \rceil$ of k (via a partial rank on S) and its limits $i_l = \mathrm{select}_a(S, (q-1) \lg^4 n)$ and $i_r = \mathrm{select}_a(S, q \lg^4 n)$. Then we see if $i_l < i \leq i_r$. In this case, we use the succinct SB-tree associated with $J_b^t[i_r] = 1$ to find the successor of i in time $\mathcal{O}(1)$. Then we use Lemma 3.1 from that position to determine in $\mathcal{O}(1)$ time if a is a τ-majority in $S[i..j]$.

[2] In that paper they find predecessors, but the problem is analogous.

If k is not in the same chunk of i, we disregard it because, in this case, there is an occurrence $S[i_l] = a$ in $S[i..j]$ that is marked in J_b^t. We will instead find separately the leftmost occurrence in $S[i..j]$ of any candidate a that is marked in $J_b^t[i..j]$, as follows. We apply Muthukrishnan's algorithm on the 1s of $J_b^t[i..j]$, to find the distinct elements of $S_b^t[\mathrm{rank}_1(J_b^t, i-1)+1, \mathrm{rank}_1(J_b^t, j)]$. Thus we obtain the leftmost sampled occurrences in $S[i..j]$ of all the τ-majorities, among other candidates. For each leftmost occurrence $S_b^t[k']$, it must be that $k = \mathrm{select}_1(J_b^t, k')$ is in the same chunk of i, and therefore we can find the successor of i using the corresponding succinct SB-tree in constant time, and then verify the candidate using Lemma 3.1.

It follows from the construction of J_b^t that the distinct elements sampled in any $S[i..j]$ must appear at least 2^t times in an interval of size $\mathcal{O}(2^b)$ containing $S[i..j]$, and so there can only be $\mathcal{O}(1/\tau)$ distinct sampled elements. Therefore, Muthukrishnan's algorithm on $J_b^t[i..j]$ gives us $\mathcal{O}(1/\tau)$ candidates to verify, in time $\mathcal{O}(1/\tau)$.

When $b < t$, we use our sequential algorithm of Section 4.2 with the only difference that, since we always find the leftmost occurrence of each candidate in $S[i..j]$, we can use Lemma 3.1 to verify the τ-majorities. Thus the algorithm uses only select and partial rank queries on S, and therefore it runs in time $\mathcal{O}(1/\tau)$ as well.

THEOREM 5.1. *Let $S[1..n]$ be a string over alphabet $[1..\sigma]$, with $\lg\sigma = \Omega(\lg\lg n)$. We can store S in $\mathcal{O}(n\lg\sigma)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

5.3. Compressed space

To reduce the space, we use the same strategy used to prove Theorem 4.5: we represent S using the optimally compressed structure of Barbay *et al.* [2014]. This time, however, closer to the original article, we use different representations for the strings S_c with alphabets of size $\sigma_c \le w$ and of size $\sigma_c > w$. For the former, we use the representation of Theorem 4.3, which uses $n_c\lg\sigma_c + o(n_c)$ bits and answers τ_c-majority queries in time $\mathcal{O}(1/\tau_c)$. For the larger alphabets, we use a slight variant of Theorem 5.1: we use the same structures G_b^t, J_b^t, C_b^t, and pointers to succinct SB-trees, except that the lower bound for b will be $\lfloor\lg(2^t\cdot g(n,\sigma))\rfloor$, for any function $g(n,\sigma) = \omega(1)$. The total space for the bitvectors G_b^t of string S_c is thus $\mathcal{O}\Big(\frac{n_c\lg\sigma_c\lg g(n,\sigma)}{g(n,\sigma)}\Big) = o(n_c\lg\sigma_c)$, whereas the other structures already used $o(n_c)$ bits (with a couple of exceptions we consider soon).

Then, representing S_c with the structure of Belazzougui & Navarro [2015, Thm. 6], so that it supports select in time $\mathcal{O}(g(n,\sigma))$ and access in time $\mathcal{O}(1)$, the total space for S_c is $n_c\lg\sigma_c + o(n_c\lg\sigma_c)$, and the whole structure uses $nH + o(n)(H+1)$ bits.

The cases where $b \ge \lfloor\lg(2^t\cdot g(n,\sigma))\rfloor$ are solved with $\mathcal{O}(1/\tau_c)$ applications of select on S, and therefore take time $\mathcal{O}((1/\tau_c)\,g(n,\sigma))$. Instead, the shorter ranges, of length $\mathcal{O}((1/\tau_c)\,g(n,\sigma))$, must be processed sequentially, as in Section 4.2. The space of the sequential algorithm can be maintained in $\mathcal{O}(1/\tau_c) = \mathcal{O}(1/\tau)$ words as follows. We cut the interval $S_c[i_c..j_c]$ into chunks of $m = \lceil 1/\tau_c\rceil$ consecutive elements, and process each chunk in turn as in Section 4.2. The difference is that we maintain an array with the τ_c-majorities a we have reported and the last position p_a we have deleted in the lists. From the second chunk onwards, we remove all the positions of the known τ_c-majorities a before processing it, $\mathrm{select}_a(S_c, \mathrm{rank}_a(S_c, p_a)+r)$, for $r = 1, 2, \ldots$; note that $\mathrm{rank}_a(S_c, p_a)$ is a partial rank query. Since select on S_c costs $\mathcal{O}(g(n,\sigma))$ and we perform $\mathcal{O}((1/\tau_c)\,g(n,\sigma))$ operations, the total time is $\mathcal{O}\big((1/\tau_c)\,g(n,\sigma)^2\big)$. Then we can retain the optimally compressed space and have any time of the form $\mathcal{O}((1/\tau)\,f(n,\sigma))$ by choosing $g(n,\sigma) = \sqrt{f(n,\sigma)}$.

There are, as anticipated, two final obstacles related to the space. The first are the $\mathcal{O}(n_c)$ bits of Muthukrishnan's structure associated with S_c to handle the case $\tau_c < 1/\sigma_c$. To reduce this space to $o(n_c)$, we sparsify the structure as in Section 3.2. The case of small τ_c is then handled in time $\mathcal{O}\big(\sigma_c\, g(n,\sigma)^2\big) = \mathcal{O}((1/\tau_c)\, f(n,\sigma))$ and the space for the sparsified structure is $\mathcal{O}(n_c/g(n,\sigma)) = o(n_c)$.

The second obstacle is the $\mathcal{O}(n_c \lg\lg n_c)$ bits used by the succinct SB-trees. Examination of the proof of Lemma 3.3 in Grossi *et al.* [2009] reveals that one can obtain $\mathcal{O}(p \lg\lg u)$ bits of space and $\mathcal{O}(\lg p/\lg\lg n)$ time if we have p elements in a universe $[1..u]$ and can store a precomputed table of size $o(n)$ that is shared among all the succinct SB-trees. We reduce the universe size as follows. We logically cut the string S_c into n_c/σ_c^2 pieces of length σ_c^2. For each symbol a we store a bitvector $B_a[1, n_c/\sigma_c^2]$ where $B_a[i] = 1$ if and only if a appears in the ith piece. These bitvectors require $\mathcal{O}(n_c/\sigma_c)$ bits in total, including support for rank and select. The succinct SB-trees are now local to the pieces: a succinct SB-tree that spans several pieces is split into several succinct SB-trees, one covering the positions in each piece. The 1s corresponding to these pieces in bitvectors B_a point to the newly created succinct SB-trees. To find the successor of position i given that it is in the same chunk of $i_r > i$, with $J_b^t[i_r] = 1$, we first compute the piece $p = \lceil i/\sigma_c^2 \rceil$ of i and the piece $p_r = \lceil i_r/\sigma_c^2 \rceil$ of i_r, and see if i and i_r are in the same piece, that is, if $p = p_r$. If so, the answer is to be found in the succinct SB-tree associated with the 1 at $J_b^t[i_r]$. Otherwise, that original structure has been split into several, and the part that covers the piece of i is associated with the 1 at $B_a[p]$. It is possible, however, that there are no elements in the piece p, that is, $B_a[p] = 0$, or that there are elements but no one is after i, that is, the succinct SB-tree associated with piece p finds no successor of i. In this case, we find the next piece that follows p where a has occurrences, $p' = \mathrm{select}_1(B_a, \mathrm{rank}_1(B_a, p) + 1)$, and if $p' < p_r$ we query the succinct SB-tree associated with $B_a[p'] = 1$ for its first element (or the successor of i). If, instead, $p' \geq p_r$, we query instead the succinct SB-tree associated with $J_b^t[i_r] = 1$, as its positions are to the left of those associated with $B_a[p'] = 1$. Therefore, successor queries still take $\mathcal{O}(1)$ time. The total number of elements stored in succinct SB-trees is still at most n_c, because no duplicate elements are stored, but now each requires only $\mathcal{O}(\lg\lg\sigma_c)$ bits, for a total space of $\mathcal{O}(n_c \lg\lg\sigma_c) = o(n_c \lg\sigma_c)$ bits. There may be up to $\sigma \cdot (n_c/\sigma^2)$ pointers to succinct SB-trees from bitvectors B_a, each requiring $\mathcal{O}(\lg n_c)$ bits, for a total of $\mathcal{O}((n_c \lg n_c)/\sigma_c) = \mathcal{O}(n_c) = o(n_c \lg\sigma_c)$, since $\sigma_c > w$.

THEOREM 5.2. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$. For any $f(n,\sigma) = \omega(1)$, we can store S in $nH + o(n)(H+1)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}((1/\tau)\, f(n,\sigma))$.*

Note that accessing a position in S still requires constant time with this representation. Further, we can obtain a version using nearly compressed space, $(1+\epsilon)nH + o(n)$ bits for any constant $\epsilon > 0$, with optimal query time, by setting $g(n,\sigma)$ to a constant value. First, use for S_c the structure of Barbay *et al.* [2014] that needs $(1 + \epsilon/3)nH + o(n)$ bits and solves access and select in constant time. Second, let κ be the constant associated with the $\mathcal{O}\Big(\frac{n_c \lg\sigma_c \lg g(n,\sigma)}{g(n,\sigma)}\Big)$ bits used by bitvectors G_b^t and the sparsified Muthukrishnan's structures. Then, choosing $g(n,\sigma) = \frac{6\kappa}{\epsilon} \lg \frac{6\kappa}{\epsilon}$ ensures that the space becomes $(\epsilon/3)n_c \lg\sigma_c$ bits, which add up to $(\epsilon/3)nH$. All the other terms of the form $o(nH)$ are smaller than another $(\epsilon/3)nH + o(n)$. Therefore the total space adds up to $(1+\epsilon)nH + o(n)$ bits. The time to sequentially solve a range of length $\mathcal{O}((1/\tau_c)\, g(n,\sigma))$ is $\mathcal{O}\big((1/\tau_c)\, g(n,\sigma)^2\big) = \mathcal{O}(1/\tau_c)$.

THEOREM 5.3. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$. For any constant $\epsilon > 0$, we can store S in $(1+\epsilon)nH + o(n)$ bits such that later, given the endpoints of a range and τ, we can return the τ-majorities for that range in time $\mathcal{O}(1/\tau)$.*

5.4. Finding range modes

While finding range modes is a much harder problem in general, we note that we can use our data structure from Theorem 5.2 to find a range mode quickly when it is actually reasonably frequent. Suppose we want to find the mode of $S[i..j]$, where it occurs occ times (we do not know occ). We perform multiple range τ-majority queries on $S[i..j]$, starting with $\tau = 1$ and repeatedly reducing it by a factor of 2 until we find at least one τ-majority. This takes time

$$\mathcal{O}\Big(\Big(1+2+4+\ldots+2^{\lceil \lg \frac{j-i+1}{\mathrm{occ}} \rceil}\Big) f(n,\sigma)\Big) = \mathcal{O}\left(\frac{(j-i+1)f(n,\sigma)}{\mathrm{occ}}\right)$$

and returns a list of $\mathcal{O}(\frac{j-i+1}{\mathrm{occ}})$ elements that includes all those that occur at least occ times in $S[i..j]$. We use rank queries to determine which of these elements is the mode. For the fastest possible time on those rank queries, we use for S the representation of Belazzougui & Navarro [2015, Thm. 8], and also set $f(n,\sigma) = \lg \lg_w \sigma$, the same time of rank. The cost is then $\mathcal{O}\Big(\frac{(j-i+1)\lg\lg_w \sigma}{\mathrm{occ}}\Big)$. The theorem holds for $\lg \sigma = \mathcal{O}(\lg w)$ too, as in this case we can use Theorem 4.5 with constant-time rank queries.

THEOREM 5.4. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H. We can store S in $nH + o(n)(H+1)$ bits such that later, given endpoints i and j, we can return the mode of $S[i..j]$ in $\mathcal{O}\Big(\frac{(j-i+1)\lg\lg_w \sigma}{\mathrm{occ}}\Big)$ time, where occ is the number of times the mode occurs in $S[i..j]$.*

6. RANGE MINORITIES REVISITED

The results obtained for range majorities can be adapted to find range minorities, which in particular improves the result of Theorem 3.2. The main idea is again that, if we test any $\lceil 1/\tau \rceil$ distinct elements, we must find a τ-minority because not all of those can occur more than $\tau(j-i+1)$ times in $S[i..j]$. Therefore, we can use mechanisms similar to those we designed to find $\mathcal{O}(1/\tau)$ distinct candidates to τ-majorities.

Let us first consider the bitvectors G_b^t defined in Section 4. We now define bitvectors I_b^t, where we flag the positions of the first 2^t and the last 2^t distinct values in each block (we may flag fewer positions if the block contains less than 2^t distinct values). Since we set $\mathcal{O}(2^t)$ bits per block, the bitvectors I_b^t use asymptotically the same space of the bitvectors G_b^t.

Given a τ-minority query, we compute b and t as in Section 4 and use rank and select to find all the 1s in the range $I_b^t[i..j]$. Those positions contain a τ-minority in $S[i..j]$ if there is one, as shown next.

LEMMA 6.1. *The positions flagged in $I_b^t[i..j]$ contain a τ-minority in $S[i..j]$, if there is one.*

PROOF. If $I_b^t[i..j]$ overlaps a block where it does not flag $2^t = \lceil 1/\tau \rceil$ distinct elements (in which case one is for sure a τ-minority), then it marks all the distinct block elements that fall inside $[i..j]$. This is obvious if $[i..j]$ fully contains the block, and it also holds if $[i..j]$ intersects a prefix or a suffix of the block, since the block marks its 2^t first and last occurrences of distinct elements. □

Just as for τ-majorities, we use I_b^t only if $1/\tau \leq \sigma$, since otherwise we can test all the alphabet elements one by one. The test proceeds using rank on S if σ is small, or using Lemma 3.1 if σ is large. We now describe precisely how we proceed.

6.1. Small alphabets

If $\lg \sigma = \mathcal{O}(\lg w)$, we use a multiary wavelet tree as in Section 4. This time, we do not run τ-majority queries on each wavelet tree node v to determine which of its children to explore, but rather we explore every child having some symbol in the range $S_v[i_v..j_v]$. To efficiently find the distinct symbols that appear the range, we store a sparsified Muthukrishnan's structure similar to the one described in Section 3.2; this time we will have no slowdown thanks to the small alphabet of S_v.

Let C_v be the array corresponding to string S_v. We cut S_v into blocks of w^β bits, and record in an array $C'_v[1..n_v/w^\beta]$ the minimum value in the corresponding block of C_v. Then, the leftmost occurrence $S[k] = p$ of each distinct symbol p in $S_v[i_v..j_v]$ has a value $C_v[k] < i_v$, and thus its corresponding block $C'_v[k']$ also holds $C'_v[k'] < i_v$. We initialize a word $E \leftarrow 0$ containing flags for the $\sigma' = w^\beta$ symbols, separated as in the final state of the word C of Section 4.3. Each time the algorithm of Muthukrishnan on C'_v gives us a new block, we apply the algorithm of Section 4.3 to count in a word C the occurrences of the distinct symbols in that block, we isolate the counters reaching the threshold $y = 1$, and compare E with E OR C. If they are equal, then we stop the recursive algorithm, since all the symbols in the range had already appeared before (see the final comments on the proof of Lemma 3.3). Otherwise, we process the subrange to the left of the block, update $E \leftarrow E$ OR C, and process the subrange to the right. When we finish, E contains all the symbols that appear in $S_v[i_v..j_v]$. In the recursive process, we also stop when we have considered $\lceil 1/\tau_v \rceil$ blocks, since each includes at least one new element and it is sufficient to explore $\lceil 1/\tau_v \rceil$ children to find a τ_v-minority (because each child contains at least one candidate). Finally, we extract the bits of E one by one as done in Section 4.3 with the use of D. For each extracted bit, we enter the corresponding child in the wavelet tree. The total time is thus $\mathcal{O}(1/\tau_v)$ and the bitvectors C_v add up to $\mathcal{O}\big(n/w^\beta\big) = o(n)$ bits in total.

The τ_u values to use in the children u of v are computed as in Section 4.3, so the analysis leading to $\mathcal{O}(1/\tau)$ total time applies. When we arrive at the leaves u of the wavelet tree, we obtain the distinct elements and compute using rank the number of times they occur in $S_u[i_u..j_u]$, so we can immediately report the first τ-minority we find.

We still have to describe how we handle the intervals that are smaller than the lower limit for b, $\lfloor 2^t \cdot w^\beta/4 \rfloor$. We do the counting exactly as in Section 4.3. We must then obtain the counters that are between 1 and $y-1$. On one hand, we use the bound $y' = 1$ and repeat their computation to obtain in $C_l \leftarrow C$ the counters that are at least 1. On the other, we compute $C \leftarrow C + (2^{2\ell} - y) \cdot (0^{k\ell+\ell-1}10^{(k-1)\ell})^{\sigma'}$ as before, and isolate the non-overflowed bits with $C_r \leftarrow (\text{NOT } C)$ AND $(0^{(k-1)\ell-1}10^{(k+1)\ell})^{\sigma'}$. Then we extract the first of the bits marked in $C \leftarrow C_l$ AND C_r and report it.

To obtain compressed space, we use the alphabet partitioning technique of Section 4.4. Once again, we must identify at most $\lceil 1/\tau \rceil$ nonempty ranges $[i_c..j_c]$ from $K[i..j]$. Those are obtained in the same way as on the multiary wavelet tree, since K is represented in that way (albeit the strings S_v are compressed). We then look for τ_c-minorities in the strings $S_c[i_c..j_c]$ one by one, until we find one or we exhaust them. The total time is $\mathcal{O}(1/\tau)$.

THEOREM 6.2. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$, with $\lg \sigma = \mathcal{O}(\lg w)$. We can store S in $nH + o(n)$ bits such that later, given the endpoints of a range and τ, we can return a τ-minority for that range (if one exists) in time $\mathcal{O}(1/\tau)$.*

Note that we can use a single representation using $nH + o(n)$ bits solving both the τ-majority queries of Theorem 4.5 and the τ-minority queries of Theorem 6.2.

6.2. Large alphabets

For large alphabets we must use Lemma 3.1 to check for τ-minorities, and thus we must find the leftmost positions in $S[i..j]$ of the τ-minority candidates. We use the same bitvectors J_b^t of Section 5, so that they store sampled positions corresponding to the 1s in I_b^t, and proceed exactly as in that section, both if $\tau < 1/\sigma$ or if $\tau \geq 1/\sigma$.

To obtain compression, we also use alphabet partitioning. We use on the multiary wavelet tree of K the method described in Section 6.1, and then complete the queries with τ_c-minority queries on the strings S_c over small or large alphabets, as required, until we find one result or exhaust all the strings. The only novelty is that we must now find τ_c-minorities sequentially for the ranges that are shorter than $\lfloor \lg(2^t \cdot g(n,\sigma)) \rfloor = \mathcal{O}((1/\tau_c)\, g(n,\sigma))$. For this, we adapt the $\mathcal{O}\big((1/\tau)\, g(n,\sigma)^2\big)$-time sequential algorithm described in Section 5.3. The only difference is that we stop as soon as we test a candidate a that turns out not to be a τ_c-majority, then reporting the τ-minority a.

Depending on whether we use Theorem 5.2 or 5.3 to represent S and how we choose $f(n,\sigma)$, we obtain Theorem 3.4 again or an improved version of Theorem 3.2:

THEOREM 6.3. *Let $S[1..n]$ be a string whose distribution of symbols has entropy H, over alphabet $[1..\sigma]$. For any constant $\epsilon > 0$, we can store S in $(1+\epsilon)nH + o(n)$ bits such that later, given the endpoints of a range and τ, we can return a τ-minority for that range (if one exists) in time $\mathcal{O}(1/\tau)$.*

In both cases, we can share the same structures to find majorities and minorities.

7. CONCLUSIONS

We have given the first linear-space data structure for parameterized range majority with query time $\mathcal{O}(1/\tau)$, even in the more difficult case of τ specified at query time. This is worst-case optimal in terms of n and τ, since the output size may be up to $1/\tau$. Moreover, we have improved the space bounds for parameterized range majority and minority, reaching in many cases optimally compressed space with respect to the entropy H of the distribution of the symbols in the sequence. While we have almost closed the problem in these terms, there are some loose ends that require further research:

— Our results for τ-majorities are worst-case time optimal, but they take $\mathcal{O}(1/\tau)$ time even if the number of majorities is $o(1/\tau)$. Is it possible to run in time $\mathcal{O}(\mathrm{occ}+1)$ when there are occ τ-majorities? Can we use $\mathcal{O}(\mathrm{occ}+1)$ instead of $\mathcal{O}(1/\tau)$ space?
— Our structure and previous ones for τ-minorities also take time $\mathcal{O}(1/\tau)$, although we are required to output only one τ-minority. Is it possible to run in $\mathcal{O}(1)$ time, or to prove a lower bound? Can we use less than $\mathcal{O}(1/\tau)$ space?
— On large alphabets, $\sigma = w^{\omega(1)}$, both for τ-majorities and τ-minorities we must use $(1+\epsilon)nH+o(n)$ bits, for any constant $\epsilon > 0$, to reach time $\mathcal{O}(1/\tau)$. With $nH+o(n)(H+1)$ bits we only have a time of the form $(1/\tau)\cdot\omega(1)$. Is it possible to close this gap?
— Our results do not improve when τ is fixed at indexing time, which is in principle an easier scenario. Is it possible to obtain better results for fixed τ?

ACKNOWLEDGMENTS

Many thanks to Patrick Nicholson for helpful comments.

REFERENCES

Barbay, J., Claude, F., Gagie, T., Navarro, G., & Nekrich, Y. 2014. Efficient fully-compressed sequence representations. *Algorithmica*, **69**(1), 232-268.

Belazzougui, D., & Navarro, G. 2014. Alphabet-independent compressed text indexing. *ACM Transactions on Algorithms*, **10**(4), article 23.

Belazzougui, D., & Navarro, G. 2015. Optimal Lower and Upper Bounds for Representing Sequences. *ACM Transactions on Algorithms*, **11**(4), article 31.

Belazzougui, D., Boldi, P., Pagh, R., & Vigna, S.. 2009. Monotone minimal perfect hashing: searching a sorted table with O (1) accesses. *Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 785–794.

Bose, P., Kranakis, E., Morin, P., & Tang, Y. 2005. Approximate range mode and range median queries. *Proc. 22nd Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 377-388.

Chan, T. M., Durocher, S., Skala, M., & Wilkinson, B. T. 2015. Linear-space data structures for range minority query in arrays. *Algorithmica*, **72**(4), 901–913.

Chan, T. M., Durocher, S., Larsen, K. G., Morrison, J., & Wilkinson, B. T. 2014. Linear-space data structures for range mode query in arrays. *Theory of Computing Systems*, **55**(4), 719–741.

Cormode, G., & Muthukrishnan, S. 2003. *Data Stream Methods*. `http://www.cs.rutgers.edu/~muthu/198-3.pdf`. Lecture 3 of Rutger's *198:671 Seminar on Processing Massive Data Sets*.

Demaine, E. D., López-Ortiz, A., & Munro, J. I. 2002. Frequency estimation of internet packet streams with limited space. *Proc. 10th European Symposium on Algorithms (ESA)*, pp. 348–360.

Durocher, S., He, M., Munro, J. I., Nicholson, P. K., & Skala, Matthew. 2013. Range majority in constant time and linear space. *Information and Computation*, **222**, 169–179.

Durocher, S., Shah, R., Skala, M., & Thankachan, S. V. 2016. Linear-Space Data Structures for Range Frequency Queries on Arrays and Trees. *Algorithmica*, **74**(1), 344-366.

Elmasry, A., Munro, J. I., & Nicholson, P. K. 2011. Dynamic range majority data structures. *Proc. 22nd International Symposium on Algorithms and Computation (ISAAC)*, pp. 150–159.

Ferragina, P., Manzini, G., Mäkinen, V., & Navarro, G. 2007. Compressed representations of sequences and full-text indexes. *ACM Transactions on Algorithms*, **3**(2), article 20.

Fischer, J. 2010. Optimal succinctness for range minimum queries. *Proc. 9th Latin American Symposium on Theoretical Informatics (LATIN)*, pp. 158–169.

Gagie, T., He, M., Munro, J. I., & Nicholson, P. K. 2011. Finding frequent elements in compressed 2D arrays and strings. *Proc. 18th Symposium on String Processing and Information Retrieval (SPIRE)*, pp. 295–300.

Golynski, A., Munro, I., & Rao, S. 2006. Rank/select operations on large alphabets: a tool for text indexing. *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 368–373.

Greve, M., Jørgensen, A. G., Larsen, K. D., & Truelsen, J. 2010. Cell probe lower bounds and approximations for range mode. *Proc. 37th International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 605–616.

Grossi, R., Orlandi, A., Raman, R., & Rao, S. S. 2009. More Haste, Less Waste: Lowering the Redundancy in Fully Indexable Dictionaries. *Proc. 26th Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 517–528.

Hon, W.-K., Shah, R., & Vitter, J. 2009. Space-Efficient Framework for Top-k String Retrieval Problems. *Proc. 50th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 713–722.

Karp, R. M., Shenker, S., & Papadimitriou, C. H. 2003. A simple algorithm for finding frequent elements in streams and bags. *ACM Transactions on Database Systems*, **28**(1), 51–55.

Karpinski, M., & Nekrich, Y. 2008. Searching for frequent colors in rectangles. *Proc. 20th Canadian Conference on Computational Geometry (CCCG)*, pp. 11–14.

Krizanc, D., Morin, P., & Smid, M. H. M. 2005. Range mode and range median queries on lists and trees. *Nordic Journal of Computing*, **12**(1), 1–17.

Lai, Y. K., Poon, C. K., & Shi, B. 2008. Approximate colored range and point enclosure queries. *Journal of Discrete Algorithms*, **6**(3), 420–432.

Misra, J., & Gries, D. 1982. Finding repeated elements. *Science of Computer Programming*, **2**(2), 143–152.

Muthukrishnan, S. 2002. Efficient algorithms for document retrieval problems. *Proc. 13th Symposium on Discrete Algorithms (SODA)*, pp. 657–666.

Navarro, G., & Thankachan, S. V. 2016. Optimal Encodings for Range Majority Queries. *Algorithmica*, **74**(3), 1082–1098.

Petersen, H. 2008. Improved bounds for range mode and range median queries. *Proc. 34th Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, pp. 418–423.

Petersen, H., & Grabowski, S. 2009. Range mode and range median queries in constant time and sub-quadratic space. *Information Processing Letters*, **109**(4), 225–228.

Pătraşcu, M. 2008. Succincter. *Proc. 49th Symposium on Foundations of Computer Science (FOCS)*, pp. 305–313.

Sadakane, K. 2007. Succinct data structures for flexible text retrieval systems. *Journal of Discrete Algorithms*, **5**(1), 12–22.

Wei, Z., & Yi, K. 2011. Beyond simple aggregates: indexing for summary queries. *Proc. 30th Symposium on Principles of Database Systems (PODS)*, pp. 117–128.